



UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
10/691,295	10/22/2003	Eric Lawrence Barsness	ROC920030239US1	2239
30206	7590	09/25/2007		
IBM CORPORATION ROCHESTER IP LAW DEPT. 917 3605 HIGHWAY 52 NORTH ROCHESTER, MN 55901-7829			EXAMINER HICKS, MICHAEL J	
			ART UNIT 2165	PAPER NUMBER
			MAIL DATE 09/25/2007	DELIVERY MODE PAPER

Please find below and/or attached an Office communication concerning this application or proceeding.

The time period for reply, if any, is set in the attached communication.

Office Action Summary	Application No. 10/691,295	Applicant(s) BARSNESS ET AL.	
	Examiner Michael J. Hicks	Art Unit 2165	

-- The MAILING DATE of this communication appears on the cover sheet with the correspondence address --

Period for Reply

A SHORTENED STATUTORY PERIOD FOR REPLY IS SET TO EXPIRE 3 MONTH(S) OR THIRTY (30) DAYS, WHICHEVER IS LONGER, FROM THE MAILING DATE OF THIS COMMUNICATION.

- Extensions of time may be available under the provisions of 37 CFR 1.136(a). In no event, however, may a reply be timely filed after SIX (6) MONTHS from the mailing date of this communication.
- If NO period for reply is specified above, the maximum statutory period will apply and will expire SIX (6) MONTHS from the mailing date of this communication.
- Failure to reply within the set or extended period for reply will, by statute, cause the application to become ABANDONED (35 U.S.C. § 133). Any reply received by the Office later than three months after the mailing date of this communication, even if timely filed, may reduce any earned patent term adjustment. See 37 CFR 1.704(b).

Status

- 1) ☒ Responsive to communication(s) filed on 25 July 2007.
- 2a) ☒ This action is **FINAL**. 2b) ☐ This action is non-final.
- 3) ☐ Since this application is in condition for allowance except for formal matters, prosecution as to the merits is closed in accordance with the practice under *Ex parte Quayle*, 1935 C.D. 11, 453 O.G. 213.

Disposition of Claims

- 4) ☒ Claim(s) 1,2,6-9,11-14,16,17 and 20 is/are pending in the application.
- 4a) Of the above claim(s) _____ is/are withdrawn from consideration.
- 5) ☐ Claim(s) _____ is/are allowed.
- 6) ☒ Claim(s) 1,2,6-9,11-14,16,17 and 20 is/are rejected.
- 7) ☐ Claim(s) _____ is/are objected to.
- 8) ☐ Claim(s) _____ are subject to restriction and/or election requirement.

Application Papers

- 9) ☐ The specification is objected to by the Examiner.
- 10) ☒ The drawing(s) filed on 22 October 2003 is/are: a) ☒ accepted or b) ☐ objected to by the Examiner.
Applicant may not request that any objection to the drawing(s) be held in abeyance. See 37 CFR 1.85(a).
Replacement drawing sheet(s) including the correction is required if the drawing(s) is objected to. See 37 CFR 1.121(d).
- 11) ☐ The oath or declaration is objected to by the Examiner. Note the attached Office Action or form PTO-152.

Priority under 35 U.S.C. § 119

- 12) ☐ Acknowledgment is made of a claim for foreign priority under 35 U.S.C. § 119(a)-(d) or (f).
- a) ☐ All b) ☐ Some * c) ☐ None of:
1. ☐ Certified copies of the priority documents have been received.
 2. ☐ Certified copies of the priority documents have been received in Application No. _____.
 3. ☐ Copies of the certified copies of the priority documents have been received in this National Stage application from the International Bureau (PCT Rule 17.2(a)).
- * See the attached detailed Office action for a list of the certified copies not received.

Attachment(s)

- | | |
|--|---|
| 1) <input type="checkbox"/> Notice of References Cited (PTO-892) | 4) <input type="checkbox"/> Interview Summary (PTO-413)
Paper No(s)/Mail Date. _____ |
| 2) <input type="checkbox"/> Notice of Draftsperson's Patent Drawing Review (PTO-948) | 5) <input type="checkbox"/> Notice of Informal Patent Application |
| 3) <input type="checkbox"/> Information Disclosure Statement(s) (PTO/SB/08)
Paper No(s)/Mail Date _____ | 6) <input type="checkbox"/> Other: _____ |

DETAILED ACTION

1. Claims 1-2, 6-9, 11-14, 16-17, and 20 Pending.
Claims 3-5, 10, 15, and 18-19 Canceled.

Response to Arguments

2. Applicant's arguments filed 7/25/2007 have been fully considered but they are not persuasive.

As per Applicants comments that Definitions 4.6, as cited in the previous action, does not exist within the disclosure of Sapia, Examiner points Applicant to indicated Section 4.2, in which Definition 4.6 appears on Page 2-6, Column 2 Though Page 2-7 Column 1. Examiner apologizes if the reference to Definition 4.6 was unclear and notes that the citation will be amended to clarify the location of Definition 4.6, but maintains that as the entirety of Section 4.2 was indicated in the citation, Definition 4.6 was included as part of the cited disclosure.

As per Applicants arguments that Sapia does not teach or suggest 'attributes', Examiner respectfully disagrees. In Section 4.2, specifically in Definition 4.6, it can be seen that the query distances are calculated on the number of transform operations that must be performed on one of the queries in order to transform it into another of the queries. Note that the number and type of these operation *deals directly with the database attributes which are specified in the queries*, as well as the operators used on

those attributes. This clearly indicates the existence of the attributes/host variables in the queries, both previous and current.

As per Applicants arguments that the query distance in Sapia is not calculated *based on attributes*, the above arguments in reference to the existence of the attributes also illustrate that Definition 4.6 indicates that the distance is based on the similarity of the attributes indicated in two queries as well as the operators used in the queries.

As per Applicants arguments that the variable in a host language is not set to a plurality of values in succession and submitted to the database, Examiner respectfully disagrees. Note that the host variable (e.g. attribute) will be written in a host language (e.g. in this case the database design language which was used) and the purpose of database attributes is to allow users to set them to values and be submitted to a database. Due to this, examiner feels that the behavior described by this particular limitation is inherent to the database attributes as disclosed by Sapia, as they must have been defined and set to values in order to be queried in the first place.

As per Applicants arguments that Sapia fails to disclose 'finding the previous statement in the history and finding the second statement following that was next in time following the previous statement in the history, Examiner respectfully disagrees. Note that the cited section of Sapia along with Definition 4.6 and Page 2-8, Column 2, Paragraph 3 of Sapia indicates that previous queries from the query log and the current

query are compared for similarity to predict next queries. Examiner asserts that even if the queries that make up the query log are from the current session, they still qualify as past queries which are in the history/log. Furthermore, note that after the previous statement is identified, it is passed to q prediction unit which then determines the next query using user pattern recognition processes (Sapia, Page 2-9, Column 2, Paragraph 5; Page 2-10, Column 1, Paragraph 1), in other words, finding the next statement in the history that the user input and predicting it as the next statement.

In light of these arguments, Examiner feels it is appropriate to maintain the current rejection.

Claim Rejections - 35 USC § 102

1. The following is a quotation of the appropriate paragraphs of 35 U.S.C. 102 that form the basis for the rejections under this section made in this Office action:

A person shall be entitled to a patent unless –

(b) the invention was patented or described in a printed publication in this or a foreign country or in public use or on sale in this country, more than one year prior to the date of application for patent in the United States.

2. Claims 1-2, 6-9, 11-14, 16-17, and 20 rejected under 35 U.S.C. 102(b) as being anticipated by Sapia.

Note that Claim 16 recites all of the limitations found in Claims 1, 6, 9, 11, and 14, and Claim 2 recites all of the limitations found in Claims 8, 12, 13, and 17.

As per Claims 1, 3, 6, 9-11, 14, and 16, Sapia discloses a server, method, apparatus, and storage device comprising: a processor; and a storage device encoded with instructions, wherein the instructions when executed on the processor comprise: finding a correlation between a first statement and a previous statement (i.e. *"If typical interaction patterns (task and/or user specific) are known, this information can be used to build query profiles for users or user groups. This information can be used to optimize the performance of an OLAP system at runtime. The idea is to use the profiles together with the know prefix of the query session to predict which queries the user is likely to ask during the rest of the session. Thus these queries or parts of the results can already be computed while the user is busy formulating his next query."* The preceding text excerpt clearly indicates that a correlation is found between queries of a present session (e.g. the first statement) and queries stored in a profile (e.g. a previous statement).) (Page 8, Column 1, Paragraph 2), wherein the previous statement is stored in a history of a plurality of statements (i.e. *"The profile information stores the user/task specific profiles (using the formalism presented in section 4). These patterns are initially constructed during the conceptual user modeling process as described earlier in this paper. Another interesting alternative is to build, validate and adapt these patterns by analyzing the query logs. This task is carried out by the profile builder which uses data mining/pattern recognition techniques to derive new profiles or adapt existing profiles."* The preceding text excerpt clearly indicates that the previous statement is stored in a history/profile.) (Page 8, Column 2, Paragraph 3), and wherein the finding the correlation further comprises finding a host variable in the previous statement in a history that matches a host variable in the first statement (i.e. See Pages 6-7, Section 4.2, specifically Definition 4.6 for examples of how the correlation (e.g. distance) between current user queries and queries in the profile are calculated using host variables (e.g. attributes). Note that attribute similarity is one of the measure determining operations to transform the queries which indicate the query distance.), wherein first data supplied for the host variable in the first statement matches previous data associated with the host variable in the

previous statement (i.e. See Pages 6-7, Section 4.2, specifically Definition 4.6 for examples of how the correlation (e.g. distance) between current user queries and queries in the profile are calculated using host variables (e.g. attributes). Note that attribute similarity is one of the measure determining operations to transform the queries which indicate the query distance.), wherein the host variable in the previous statement and the host variable in the first statement comprise a variable in a host language that is set to a plurality of values in succession and submitted to a database (i.e. See Pages 6-7, Section 4.2, specifically Definition 4.6 for examples of how the correlation (e.g. distance) between current user queries and queries in the profile are calculated using host variables (e.g. attributes). Note that the attributes are written in a host language (e.g. the language used to design the database).), predicting a second statement based on the previous statement (i.e. *"The core of the architecture is a prediction unit, which is located between the query processor and the frontend tool. It has the same interface to the user interface as the query processor (e.g. MDX) and predicts the possible next queries of the user from the status of the current session (the queries asked so far) and the profile information. Following a predefined strategy it passed predicted queries to the query processor for speculative execution. The results are stored in a speculative result cache."* The preceding text excerpt clearly indicates that a second statement (e.g. next possible query) is predicted based on the first statement (e.g. a query which was asked during the session).) (Page 8, Column 2, Paragraph 3; Page 9, Column 1, Paragraph 1), wherein the predicting further comprises finding the second statement that was next in time following the previous statement in the history (i.e. *"If not, the query is passed to the query scheduler which is responsible for passing the query on to the OLAP query processor. In any case the session state for the current session is updated (i.e. the current query prototype is added to the session prefix). This information together with the profile information is used by the 'prediction model'-process which generates queries that are most likely to be asked next and estimates their probability. In this process it makes use of the distance measure defined in section 4 as an approximation of similarity."* The preceding text excerpt clearly indicates that the predicted

statement/query is next in time following the previous statement in the history/profile. Note the description of the distance measure (Definition 4.6.) (Page 9, Column 1, Paragraph 2), wherein the previous statement and the second statement comprise commands that were previously executed against a database (i.e. *"The profile information stores the user/task specific profiles (using the formalism presented in section 4). These patterns are initially constructed during the conceptual user modeling process as described earlier in this paper. Another interesting alternative is to build, validate and adapt these patterns by analyzing the query logs. This task is carried out by the profile builder which uses data mining/pattern recognition techniques to derive new profiles or adapt existing profiles."* The preceding text excerpt clearly indicates that the previous statement and the second statement were both queries/commands that were previously executed in against the database, either by a user (e.g. query log method) or during the user modeling process (see Section 4.) (Page 8, Column 2, Paragraph 3), executing the first statement against the database (i.e. *"A cost model process passes the estimated database execution costs of the queries to a decision model process which decides if the query should be executed speculatively. In this case it passes the query to the query scheduler with a flag that the results of this query should be stored in the cache instead of being sent back to the user immediately. When the query scheduler receives a new query for execution, it first checks if such a query is already being executed at the moment. In this case the query is not passed on but answered using the results of the running query. This case can occur if a query that is being speculatively executed is actually asked by the user before the execution is finished. The results of all queries are either passed to the user or stored in the speculative result cache. The cache uses an appropriate replacement strategy (e.g. a time stamp method)."* The preceding text excerpt clearly indicates that the first statement (e.g. a statement which is not being speculatively executed) will be executed against the database and the results returned to the user.) (Page 9, Column 1, Paragraph 3, Column 2, Paragraph 1), retrieving at least one page from a database based on the second statement (i.e. *"A cost model process passes the estimated database execution costs of the queries to a decision model process which decides if the query should be executed speculatively. In this case it passes the query to the query scheduler with*

a flag that the results of this query should be stored in the cache instead of being sent back to the user immediately. When the query scheduler receives a new query for execution, it first checks if such a query is already being executed at the moment. In this case the query is not passed on but answered using the results of the running query. This case can occur if a query that is being speculatively executed is actually asked by the user before the execution is finished. The results of all queries are either passed to the user or stored in the speculative result cache. The cache uses an appropriate replacement strategy (e.g. a time stamp method)." The preceding text excerpt clearly indicates that the second statement (e.g. a statement which is being speculatively executed) will be executed against the database and the results (e.g. the at least one page) will be stored in a cache.) (Page 9, Column 1, Paragraph 3, Column 2, Paragraph 1), wherein the retrieving further comprises executing the second statement against the database (i.e. "A cost model process passes the estimated database execution costs of the queries to a decision model process which decides if the query should be executed speculatively. In this case it passes the query to the query scheduler with a flag that the results of this query should be stored in the cache instead of being sent back to the user immediately. When the query scheduler receives a new query for execution, it first checks if such a query is already being executed at the moment. In this case the query is not passed on but answered using the results of the running query. This case can occur if a query that is being speculatively executed is actually asked by the user before the execution is finished. The results of all queries are either passed to the user or stored in the speculative result cache. The cache uses an appropriate replacement strategy (e.g. a time stamp method)." The preceding text excerpt clearly indicates that the second statement (e.g. a statement which is being speculatively executed) will be executed against the database and the results (e.g. the at least one page) will be stored in a cache.) (Page 9, Column 1, Paragraph 3, Column 2, Paragraph 1), storing the at least one page in a cache (i.e. "A cost model process passes the estimated database execution costs of the queries to a decision model process which decides if the query should be executed speculatively. In this case it passes the query to the query scheduler with a flag that the results of this query should be stored in the cache instead of being sent back to the user immediately. When the query

scheduler receives a new query for execution, it first checks if such a query is already being executed at the moment. In this case the query is not passed on but answered using the results of the running query. This case can occur if a query that is being speculatively executed is actually asked by the user before the execution is finished. The results of all queries are either passed to the user or stored in the speculative result cache. The cache uses an appropriate replacement strategy (e.g. a time stamp method)." The preceding text excerpt clearly indicates that the second statement (e.g. a statement which is being speculatively executed) will be executed and the results (e.g. the at least one page) will be stored in a cache.) (Page 9, Column 1, Paragraph 3, Column 2, Paragraph 1), and executing a next statement against the at least one page in the cache (i.e. *"In figure 8 the dataflow and the processes inside the prediction unit are shown (the parts where the formalism presented in this paper is used are marked gray). The frontend passes a new query to the request handler. The handler checks if this query can be answered from the speculative result cache, i.e. has been correctly predicted."* The preceding text excerpt clearly indicates that a next statement (e.g. a query which is entered after the first statement, and after the result of the second statement has been placed in the cache) is executed and a check is made to see if the result of the next statement exists in the cache. As such, the next statement is executed against the page (e.g. result) in the cache.) (Page 9, Column 1, Paragraph 2), wherein the next statement follows the first statement in time (i.e. *"In figure 8 the dataflow and the processes inside the prediction unit are shown (the parts where the formalism presented in this paper is used are marked gray). The frontend passes a new query to the request handler. The handler checks if this query can be answered from the speculative result cache, i.e. has been correctly predicted."* The preceding text excerpt clearly indicates that a next statement (e.g. a query which is entered after the first statement, and after the result of the second statement has been placed in the cache) is executed and a check is made to see if the result of the next statement exists in the cache.) (Page 9, Column 1, Paragraph 2), and wherein the host variable in the next statement matches the host variable in the second statement (i.e. *"In figure 8 the dataflow and the processes inside the prediction unit are shown (the*

parts where the formalism presented in this paper is used are marked gray). The frontend passes a new query to the request handler. The handler checks if this query can be answered from the speculative result cache, i.e. has been correctly predicted." The preceding text excerpt clearly indicates that a next statement (e.g. a query which is entered after the first statement, and after the result of the second statement has been placed in the cache) is executed and a check is made to see if the result of the next statement exists in the cache. Note that in order to have the same result sets (e.g. in order for the result of the next statement to exist in the speculative cache) the next statement and the second statement must share host variables.) (Page 9, Column 1, Paragraph 2).

As per Claims 2, 8, 12, 13, and 17, Sapia discloses the retrieving further comprises: retrieving the at least one page asynchronously from executing the first statement against the database and storing the at least one page in a cache (i.e. *"The core of the architecture is a prediction unit, which is located between the query processor and the frontend tool. It has the same interface to the user interface as the query processor (e.g. MDX) and predicts the possible next queries of the user from the status of the current session (the queries asked so far) and the profile information. Following a predefined strategy it passed predicted queries to the query processor for speculative execution. The results are stored in a speculative result cache...A cost model process passes the estimated database execution costs of the queries to a decision model process which decides if the query should be executed speculatively. In this case it passes the query to the query scheduler with a flag that the results of this query should be stored in the cache instead of being sent back to the user immediately. When the query scheduler receives a new query for execution, it first checks if such a query is already being executed at the moment. In this case the query is not passed on but answered using the results of the running query. This case can occur if a query that is being speculatively executed is actually asked by the user before the execution is finished. The results of all queries are either passed to the user or stored in the speculative result cache. The cache uses an appropriate replacement strategy (e.g. a time stamp method)." The preceding text excerpt clearly*

indicates that the one page (e.g. the result set) is retrieved asynchronously from the executing of the first statement and that the one page (e.g. result set) is stores in a speculative result cache.) (Page 8, Column 2, Paragraph 3; Page 9, Column 1, Paragraphs 1, 3, Column 2, Paragraph 1).

As per Claim 7, Sapia discloses means for saving the first statement in the history (i.e. *"Another interesting alternative is to build, validate and adapt these patterns by analyzing the query logs. This task is carried out by the profile builder which uses data mining/pattern recognition techniques to derive new profiles or adapt existing profiles...Another interesting point is the derivation of user/task profiles from the saved interaction logs using data mining and pattern recognition techniques."* The preceding text excerpt clearly indicates that the queries which are used in the session may be saved to query logs and later inserted into the profiles/history.) (Page 8, Column 2, Paragraph 3; Page 9, Column 2, Paragraph 4; Page 10, Column 1, Paragraph 1).

As per Claim 20, Sapia discloses the finding the correlation further comprises: finding the previous statement, wherein the previous statement is associated with a same job as the first statement (i.e. *"If typical interaction patterns (task and/or user specific) are known, this information can be used to build query profiles for users or user groups. This information can be used to optimize the performance of an OLAP system at runtime. The idea is to use the profiles together with the know prefix of the query session to predict which queries the user is likely to ask during the rest of the session. Thus these queries or parts of the results can already be computed while the user is busy formulating his next query."* The preceding text excerpt clearly indicates that the previous statement (e.g. the query found in the user profile) and the first statement (e.g. the known prefix f the query session) are associated with the same job (e.g. are used in the similar query sessions or are task specific).) (Page 8, Column 1, Paragraph 2).

THIS ACTION IS MADE FINAL. Applicant is reminded of the extension of time policy as set forth in 37 CFR 1.136(a).

A shortened statutory period for reply to this final action is set to expire THREE MONTHS from the mailing date of this action. In the event a first reply is filed within TWO MONTHS of the mailing date of this final action and the advisory action is not mailed until after the end of the THREE-MONTH shortened statutory period, then the shortened statutory period will expire on the date the advisory action is mailed, and any extension fee pursuant to 37 CFR 1.136(a) will be calculated from the mailing date of the advisory action. In no event, however, will the statutory period for reply expire later than SIX MONTHS from the mailing date of this final action.

Points of Contact

Any inquiry concerning this communication or earlier communications from the examiner should be directed to Michael J. Hicks whose telephone number is (571) 272-2670. The examiner can normally be reached on Monday - Friday 10:00a - 7:00p.

If attempts to reach the examiner by telephone are unsuccessful, the examiner's supervisor, Jeffrey Gaffin can be reached on (571) 272-4146. The fax phone number for the organization where this application or proceeding is assigned is 571-273-8300.

Art Unit: 2165

Information regarding the status of an application may be obtained from the Patent Application Information Retrieval (PAIR) system. Status information for published applications may be obtained from either Private PAIR or Public PAIR. Status information for unpublished applications is available through Private PAIR only. For more information about the PAIR system, see <http://pair-direct.uspto.gov>. Should you have questions on access to the Private PAIR system, contact the Electronic Business Center (EBC) at 866-217-9197 (toll-free). If you would like assistance from a USPTO Customer Service Representative or access to the automated information system, call 800-786-9199 (IN USA OR CANADA) or 571-272-1000.

Michael J Hicks
Art Unit 2165
Phone: (571) 272-2670
Fax: (571) 273-2670

Muhammad Akel Jaleel
[Signature]
Primary Examiner
9/20/07